Recovery and Common Sense Reasoning

Principles of Nonmonotonic Theory Recovery

Cees Witteveen
Delft University of Technology,
Department of Mathematics and Computer Science
P.O.Box 356, 2600 AJ Delft, The Netherlands.
witt@cs.tudelft.nl

Abstract

We use common-sense reasoning to make predictions about what normally will be the case. Such reasoning is captured by using a nonmonotonic semantics selecting a set of intended models from a larger collection of models. Such a selection process, however, might fail: although a theory might have a nonempty set of non-intended models, the set of intended models might be empty. We call such a theory nonmonotonically inconsistent. In order to restore consistency, we want to recover the theory, replacing it by a closely related one that has intended models.

In this paper we present some rationality postulates for recovery of common-sense theories. As a result, we show that, contrary to intuition, in most cases contractions of a common-sense theory are inadequate to restore consistency, and that the only minimal recovery operations are expansions of it.

Keywords: Recovery, Common Sense Reasoning, Nonmonotonic Logic.

1 Introduction

Suppose you believe that (i) normally, adults can read and (ii) normally, researchers read research reports. Furthermore, it happens that you hear that (iii) Fred is an adult. Being a rational person, you guess that (g1) Fred can read. Then someone tells you that, actually, (iv) Fred is a researcher. This offers no problem to you: now you also believe that (g2) he will read research reports. Suppose that later on you learn that in fact (v) Fred does not read a single research report. Then you have to give up (g2) and instead you will guess that, after all, (g2') Fred is not a normal researcher. At the same time, however, it is perfectly reasonable to maintain the previous conclusion (g1) that Fred can read.

Common sense theories like the one above are used to derive conclusions about what will be the case if the world is behaving normally. The machinery we use to derive such conclusions (or *intelligent guesses*) from these theories is provided by *nonmonotonic logics* such as default theory, auto-epistemic logic or nonmonotonic logic programming. And, indeed, they are perfectly able to derive (g1) from (i) - (iii) and the conclusions (g1) and (g2) from (i) - (iv). It may come as a surprise, however, that they are not able to derive (g1) and (g2') as a rational conclusion from (i) - (v). The reason is that these nonmonotonic logics are not able to find a suitable (*intended*) model, since, expecting the world to be a normal one, no intended model can be found where besides (i) - (iv), (v) also holds. So according to these logics, the theory consisting of the statements (i) - (v) is *non-monotonically inconsistent* (nm-inconsistent) and does not make sense if the world is as normal as we expected it to be.

Theory Recovery

The problem how to interpret a theory that, according to some semantics Sem, does not have an acceptable meaning, constitutes a *theory recovery problem*: given a Sem-inconsistent theory T, how do we construct a reasonable alternative T' to T such that (i) T' is Sem-consistent and (ii) the information content of T' resembles T as much as possible.

If Sem is a classical semantics, there is a well-known framework for theory recovery: the AGM-framework for theory-revision (see [2]). Since a classically inconsistent theory does not have any model at all, it comes as no surprise that the proposed recovery operation is a *contraction* of the original theory T. That is, a reasonable alternative T' is obtained by *removing* some information from T. The AGM-rationality postulates for recovery ensure that the information loss by contracting T is minimized with respect to an underlying preference relation, the entrenchment relation. This relation is used to capture the commitment of a rational agent to the information contained in the theory: roughly the idea is that the deeper some information is entrenched in the belief structure of the agent, the more reluctant the agent will be to give it up.

Recovery and nonmonotonic semantics

If a nonmonotonic semantics is used, it is not evident that the AGM-approach can be applied. For, using a nonmonotic semantics means that only a subset of some larger collection of models of the theory is selected as constituting the preferred or *intended* models of the theory. Such a nonmonotonic semantics, however, might fail to provide intended models for a common sense theory T, even if T itself has other (non-intended) models. In this case, T is called *nonmonotonically inconsistent* (nm-inconsistent). For example, the theory consisting of the statements (i)- (v) is inconsistent under a default-theory interpretation.

For an nm-inconsistent theory it may happen that both a theory T' containing T and a theory T'' contained in T have intended models[1]. Therefore, in contrast to recovery of classical theories, not only contraction, but also *expansion* or a *mixture* of both operations, in principle, might be used to recover a nonmonotonic theory[2].

It comes as no surprise that in the existing literature of nonmonotonic theory recovery, one can find proposals for *theory contraction* (e.g. [5]), for *theory expansion* (see [1, 6]) as well as a *mixture of expansion and contraction* to restore the consistency of such a theory (e.g. [3]).

In all these proposals however, we do not find a justification for using a particular recovery operation. Therefore, the goal of this paper is to give a *principled approach to theory recovery* by offering a framework for recovery with nonstandard semantics. We will show that, in particular for nonmonotonic theory recovery, only recovery by expansion is suitable as a recovery operation. Due to lack of space, all proofs of results have been omitted (but, see [8] for proofs of results).

Our approach

With an intended nonmonotonic semantics *Int*

[1]For example, if T is the nm-inconsistent theory consisting of the statements (i) - (v), the theory T' resulting from adding the fact that (vi) Fred is not a normal researcher will have intended models under a default semantics. And also does the theory T'' consisting of (i) - (iv).

[2]This also implies that the notion of an entrenchment relation may not be a useful concept to use in nonmonotonic theories, since it is prepared for recovery by contraction strategies

to reason about a theory T, a subset of models, the *intended models*, is selected from some given collection of models for T. These intended models are useful for making predictions about a normal world. Sometimes, however, the world is not as normal as we think it is. In such apparently abnormal circumstances we should not be surprised to find out that intended models of a theory T cannot be used to give a satisfactory description. So we do not consider *Int* to be defective: it is perfectly possible that there exist intended models for T *modulo* the abnormalities detected. That is, if somehow we take into account these apparent abnormalities, we might be able to use our intended semantics. To be able to do so, we will try to represent such abnormalities *syntactically*, by modifying the current theory T. We therefore adhere to the principle *change the theory and not the logic to reason with.* So what we are looking for is a recovery T' of T, such that T' does have intended models and resembles T. It should resemble T in that these intended models are the intended models of the original theory T *modulo* the abnormalities observed. But how do we obtain such a theory T'? To this end we use a class *BCK* of *background semantics*. This class is used to fall back on weaker[3] interpretations of the theory when *Int* fails to provide an interpretation. Using a semantics $S \in BCK$ we are able to select some non-intended models of T that take into account some abnormalities. Of course we want to include as few abnormalities as seems to be necessary. Therefore, we prefer those semantics S in *BCK* that come as close to the intended semantics *Int* as possible and, moreover, are able to provide a meaning to T. Let us call such a semantics a *maximal background semantics*. Since a maximal background semantics S is weaker than *Int*, it allows for more abnormalities to occur and may give us some clues about which abnormalities we have to be prepared for. We may consider the models selected by such an S as a first *approximation* of the set of intended models of the theory modulo the apparent abnormalities. So, given some *Int*-inconsistent theory T, what about theories T' that have the same maximal background models as T has? It seems to be reasonable to consider them as theories closely resembling T. Among such theories there might be a theory T' such that T' does have intended models. Such a theory T' might be considered as a syntactical reformulation of T that comes as close to its original meaning as possible.

Hence, we decide to recover the original theory T by applying a *transformation* R to T such that (i) $R(T)$ is *Int*-consistent and (ii) T and $R(T)$ are indistinguishable under *BCK*. This guarantees that we do not allow for more abnormalities to occur than is necessary.

Remark. Notice that the use of a maximal background semantics also guarantees the existence of a lower bound on the information loss associated with the recovery process: since the intended models of $R(T)$ are included in the background models of $R(T)$ and these are equal to the background models of T, we keep at least as much information as can be obtained from the background semantics.

Organisation of the paper

We will develop a formal framework for recovery based on these intuitive ideas. Before we present this framework, we first discuss some necessary preliminaries, needed to understand the formal details. Our main question is: given this framework, which type of recovery operations are adequate to deal with nonmonotonic theory recovery. By giving a general characterization of nonmonotonic theories we are able to show that under very general conditions *contraction* can be excluded and, if we want to change the theory as little as possible, we should *expand* our theory.

[3] For they allow more abnormalities to occur and therefore can predict less.

2 Preliminaries

Given a language L, a *theory* T is any subset of L. We assume to have a way to assign to each T some (possibly empty) set of models $Mod(T)$ in some specified class. For any class of theories $\mathcal{T}$ a *semantics* Sem then is a way to associate consequences φ to some $T \in \mathcal{T}$, based on $Mod(T)$. Such a semantics is called *well-behaved* w.r.t. T if $Sem(T)$ is defined and is not equal to L. We often identify a semantics Sem with a *consequence operation* $C_{Sem} : 2^L \rightarrow 2^L$, where the set $C_{Sem}(T)$ denotes the set of valid conclusions obtainable from T using the semantics Sem. In this paper, we stipulate that $C_{Sem}(T) = L$ in case $Sem(T)$ is not defined. Generalizing the above, we say that a consequence operator C is well-behaved w.r.t. T, if $C(T) \neq L$. We focus theories that have more than one semantics. Besides our intended semantics *Int* characterized by an associated consequence operator C_{int}, we consider a class of background semantics *BCK* for C_{int}. The class *BCK* is characterized by a corresponding class $\mathcal{C}_{BCK}$ of consequence operators C_{bck} for every $bck \in BCK$. The class $\mathcal{C}_{BCK}$ contains consequence operators that are weaker than C_{int}: we say that C_{int} satisfies the following property of *supra-inferentiality*: For all $A \subseteq L$ and $C_{bck} \in \mathcal{C}_{BCK}$, $C_{bck}(A) \subseteq C_{int}(A)$.

A *recovery framework* is a tuple $\mathcal{R} = (\mathcal{T}, C_{int}, \mathcal{C}_{BCK}, R)$, where $\mathcal{T}$ is a set of theories, C_{int} is the intended semantics for every $T \in \mathcal{T}$, $\mathcal{C}_{BCK}$ the background semantics for *Int* and R is a computable function $R : \mathcal{T} \rightarrow \mathcal{T}$ called the *recovery operator*. Finally, given a $T \in \mathcal{T}$, we call $C_{bck} \in \mathcal{C}_{BCK}$ a $\mathcal{C}_{BCK}$*-maximal well-behaved* consequence operator iff for every $C \in \mathcal{C}_{BCK}$ such that $C(T) \neq L$, $C_{bck}(T) \subseteq C(T)$ implies that $C_{bck}(T) = C(T)$. That is, a $\mathcal{C}_{BCK}$-maximal well-behaved consequence operator w.r.t. a theory T is a most informative consequence operator in $\mathcal{C}_{BCK}$ w.r.t. *Int*

We want to state some general results about the properties a suitable recovery operator should have. These properties partly depend on some abstract properties of the consequence operators used. Therefore we recall (see e.g. [4]) some general properties along which one can classify consequence operators: $A \subseteq C(A)$ (*Inclusion*), $C(A) = C(C(A))$ (*Idempotency*), $A \subseteq B$ implies $C(A) \subseteq C(B)$ (*Monotony*), $A \subseteq B \subseteq C(A)$ implies $C(B) \subseteq C(A)$ (*Cut*) and $A \subseteq B \subseteq C(A)$ implies $C(A) \subseteq C(B)$ (*Cautious Monotony*).

A *classical* inference operation C will also be denoted by Cn. An inference operation C is called *Tarskian*[4] if it satisfies *Inclusion*, *Idempotency* and *Monotony*, and it satisfies *Cumulativity* if both *Cut* and *Cautious Monotony* hold for C. Finally, C is called a *cumulative* inference operation, if it satisfies *Inclusion* and *Cumulativity*.

As we have shown in [7, 8], we need the following weaker variants of *Cut* and *Cautious Monotony* to distinguish non-monotonic logics from other non-standard logics:

$A \subseteq B \subseteq C(A)$ and $C(A) \neq L$ implies $C(B) \neq L$ (*Weak Cut*)

and

$A \subseteq B \subseteq C(A)$ and $C(B) \neq L$ implies $C(A) \neq L$ (*Weak Monotony*).

Slightly abusing language, we say that a class of consequence operators $\mathcal{C}$ satisfies a property P if every element of $\mathcal{C}$ satisfies P.

In [8] we have shown that all mainstream non-monotonic semantics used to model commonsense reasoning patterns[5] satisfy *Weak Cut* and none of them satisfies *Weak Monotony*. This result will be used to characterize suitable recovery operators for nonmonotonic logics.

[4]In particular, the classical consequence operator Cn is a Tarskian consequence operator.

[5]In fact we could prove that this result holds irrespective from the mode in which the inference operators based on them are used.

3. Rationality postulates for recovery

Let $\mathcal{R} = (\mathcal{T}, C_{int}, \mathcal{C}_{BCK}, R)$ be a recovery framework. Based upon the ideas developed in the introduction we will formulate the following rationality postulates that constrain the choice of a suitable recovery function given an intended semantics C_{int}, a class of background semantics $\mathcal{C}_{BCK}$ and a theory $T \in \mathcal{T}$.

R1. Success: $C_{int}(R(T)) \neq L$ whenever there exists a background semantics $C_{bck} \in \mathcal{C}_{BCK}$ such that $C_{bck}(T) \neq L$. This means that the recovery should be *successful*: if in the background semantics, one can attach a meaning to T, $R(T)$ should be well-behaved with respect to the intended semantics.

R2. Conservativity: $R(T) = T$ whenever $C_{int}(T) \neq L$.
This postulate guarantees that recovery is done in a *conservative* way: a recovery only leads to a change of T if it is necessary to do so, i.e, if one is unable to assign T a meaning under the intended semantics.

R3. Background preservation: Whenever C_{bck} is $\mathcal{C}_{BCK}$-maximal well-behaved w.r.t. T, $C_{bck}(T) \subseteq C_{bck}(R(T))$. Since the original theory is meaningful under the background semantics, we do not want to lose information obtainable from the original theory when using the transformed theory $R(T)$.

R4. Background inclusion: Whenever C_{bck} is $\mathcal{C}_{BCK}$-maximal well-behaved w.r.t. T, $C_{bck}(R(T)) \subseteq C_{bck}(T)$. This postulate constrains the recovery R by requiring that all (maximal well-behaved) background consequences of $R(T)$ should be derivable from $C_{bck}(T)$. That is, we are not allowed to add some new information when using the transformed theory.

The intention of these postulates is to characterize recovery operations that are both intuitively acceptable and *successful*:

Definition 1 We say that a recovery framework $\mathcal{R} = (\mathcal{T}, C_{int}, \mathcal{C}_{BCK}, R)$ is successful if for every $T \in \mathcal{T}$, $R(T)$ satisfies the postulates R1 to R4.

This does not exclude recovery frameworks that are successful in a *trivial* way, for example if $C_{bck}(T)$ is not well-behaved for any $T \in \mathcal{T}$ and $C_{bck} \in \mathcal{C}_{BCK}$, or $C_{int}(T)$ is well-behaved for every $T \in \mathcal{T}$. Therefore, we define a *non-trivially* successful recovery framework as follows:

Definition 2 Let $\mathcal{R} = (\mathcal{T}, C_{int}, \mathcal{C}_{BCK}, R)$ be a successful recovery framework. We say that $\mathcal{R}$ is *non-trivially* successful if there exists at least one $T \in \mathcal{T}$ such that there exists a $C_{bck} \in \mathcal{C}_{BCK}$ with $C_{bck}(T)$ is well-behaved and $C_{int}(T)$ is not well-behaved.

How to recover

The recovery postulates R1-R4 restrict the background semantics to a most informative semantics weaker than the intended semantics and the class of possible recovery operations to the ones that are considered to be *acceptable*, i.e., well-behaved and invariant under the background semantics.

Note, that in general, we do not know anything about the class of background semantics, except that (by the property of *Supra-inferentiality*) they are weaker than the intended semantics. Now suppose that we know that our intended semantics is nonmonotonic and satisfies *Weak Cut* (but not *Weak Monotony*). Could we say something about the identity of allowable recovery operations in successful nonmonotonic recovery functions? The answer is yes, we can. Even if the only thing we know is that the background semantics all satisfy *Inclusion* (a very weak assumpwe know that someone (Fred) is an adult. Then the default theory $\Delta_1 = (W_1, D)$ has exactly one (intuitively correct) extension and $C_{int}(W_1) = Cn(\{adult, can\_read\})$ as we expect. If we learn that this person is a researcher, i.e. $W_2 = \{adult, researcher\}$, we obtain the default theory $\Delta_2 = (W_2, D)$ where, as we expect, $C_{int}(W_2) = Cn(\{adult, researcher, can\_read, read\_reports\})$. But, if $W_3 = \{adult, researcher, \neg read\_reports\}$, $\Delta_3 = (W_3, D)$ does not have a Reiter extension, although it seems intuitively right to conclude that $ab\_researcher$ will hold. In our framework, we are able to recover the theory W_3 in such a way that $ab\_researcher$ will be in an extension of the recovered theory. First, let us look for a suitable background semantics. We choose the minimal extension semantics[8], for which Δ_3 has an extension. Under this semantics, the unique minimal extension for W_3 is $E_{min} = Cn(\{adult, researcher, \neg read\_reports, ab\_researcher\})$. Taking this semantics as our background semantics, we have $C_{bck}(W_3) = E_{min}$. Since we know that the consequence operator C_{int} based on the Reiter semantics satisfies *Weak Cut* and the minimal extension semantics is cumulative, according to Theorem 4 we cannot retract information from W_3 to obtain a Reiter extension for Δ_3. In fact, if we want to change W_3 in a minimal way, we are forced to apply an expansion. We decide to apply the expansion: $R(W_3) = W_3 \cup \{ab\_researcher\}$ and, indeed, we observe that (i) $C_{bck}(W_3) = C_{bck}(R(W_3))$ while (ii) $C_{int}(R(W_3))$ has a unique extension equal to E_{min}, so for this theory, the recovery operator satisfies the postulates.

remember that, actually, in Default Theory the defaults act as inference rules. The set W constitutes your observations.

[8]A minimal extension of a default theory $\Delta = (W, D)$ is a minimal set E containing W and closed under classical consequences and application of default rules.

Conclusion

We have presented a framework and some postulates for recovery of nonmonotonic theories. We have shown that in case the intended semantics is a mainstream nonmonotonic semantics, under very general conditions set for the background semantics, recovery cannot be accomplished by contraction operations. This distinguishes nonmonotonic recovery from the AGM framework for recovery of classical theories. This leaves only room for expansions and mixed recovery operations. As a special case, the Contradiction Removal framework developed by Pereira and Alferes (see [1]), satisfies our first three rationality postulates and makes use of expansions as recovery operators. Taking a classical semantics as the background semantics, the Contradiction Removal Semantics is a special recovery framework in which the postulates R1-R3 are satisfied. This means that (i) it cannot be applied if the intended semantics satisfies *Weak Cumulativity* and (ii) since the Contradiction Removal Semantics aims at adding a minimal set of revisions, the expansion approach can be justified by pointing out that contraction never can be an option. Our results show that, whenever R is a mixed recovery that satisfies the postulates R1-R4, it can always be replaced by a successful expansion that does not produce more changes. In particular, we have shown that whenever the background semantics is cumulative, syntactically minimal recovery operators for nonmonotonic theories have to be expansions in order to be successful. This result can be related to the mixed recovery approach of Inoue and Sakama (see [3]). Our results show that, whenever R is a mixed recovery that satisfies the postulates R1-R4, it can always be replaced by a successful expansion that does not produce more changes. Finally, in a case study of recovery in nonmonotonic logic programming[8], we have shown that a stable model for a classical consistent program tion, indeed) this already is sufficient to exclude contractions as recovery operators:

Definition 3 A recovery function R is called an *expansion* if for all $T \in \mathcal{T}$ we have $T \subseteq R(T)$ and R is called a *contraction* if for all $T \in T$ we have $R(T) \subseteq T$.

Theorem 4 Let $\mathcal{R} = (\mathcal{T}, C_{int}, \mathcal{C}_{BCK}, R)$ be a recovery framework, where $\mathcal{C}_{BCK}$ satisfies *Inclusion*, C_{int} satisfies *Weak Cut* and R is a contraction. Then $\mathcal{R}$ cannot be nontrivially successful w.r.t. the postulates R1-R4.

So we are justified to conclude that *in general, contractions are not useful in non-monotonic theory recovery.*

Hence, we do not agree with approaches as [5]. How about mixed recovery operations where some information is added and other information is removed from the theory? This approach is advocated in [3] to recover from non-monotonic inconsistency. The authors propose to revise a theory T by means of a *minimal set* of additions I and removals O such that $R(T) = T + I - O$ has an acceptable model. Using our framework, we have two objections against their approach. First of all, we can easily show that, under very mild conditions, expansions can be used to represent other more general recovery functions whenever these are successfully applicable.

Theorem 5 Let C_{int} satisfy *Weak Cut* and suppose that $\mathcal{C}_{BCK}$ is *cumulative*. Then there exists a successful recovery framework $\mathcal{R} = (\mathcal{T}, C_{int}, \mathcal{C}_{BCK}, R)$ satisfying the postulates R1-R4 iff there exists a successful recovery framework $\mathcal{R}' = (\mathcal{T}, C_{int}, C_{bck}, R')$, where R' is an expansion.

Theorem 5 shows that using a mainstream non-monotonic logic and a cumulative background semantics, expansions are able to characterize successful recovery frameworks.

Secondly, whenever one wants to recover a theory T by changing it in a minimal way, one *has* to use expansions. Let us define a recovery framework a *minimal change* recovery framework if the recovery operator R minimizes the difference between T and $R(T)$:

Definition 6 We call $\mathcal{R} = (\mathcal{T}, C_{int}, \mathcal{C}_{BCK}, R)$ a successful *minimal change recovery framework* if for every successful recovery framework $\mathcal{R} = (\mathcal{T}, C_{int}, \mathcal{C}_{BCK}, R')$ and every $T \in \mathcal{T}$ it holds that $R(T) \ominus T \subseteq R'(T) \ominus T$[6].

It is not difficult to see that the only recovery operators that can be used in a successful minimal change recovery framework are expansions, if we use a cumulative background semantics and an intended semantics satisfying *Weak Cut.*

Theorem 7 Let $\mathcal{R} = (\mathcal{T}, C_{int}, \mathcal{C}_{BCK}, R)$ be a nontrivial successful minimal change recovery framework where $\mathcal{C}_{BCK}$ is *cumulative* and C_{int} satisfies *Weak Cut.* Then R has to be an expansion.

So in case of *minimal-change* recovery, expansions are the only successful recovery functions and mixed recovery operators can be abandoned.

Example 8 For readers acquainted with Reiter's default theory, let us model the example given in the introduction with Reiter's default theory.

Consider the following set of defaults: $D = \{\frac{adult;\neg ab\_adult}{can\_read}, \frac{researcher;\neg ab\_researcher}{read\_reports}\}$. We use this set of defaults to construct a consequence operator C_{int} that, given a set of sentences W returns the intersection $C_{int}(W)$ of all Reiter extensions of the default theory $\Delta = (W, D)$[7]. Suppose $W_1 = \{adult\}$, i.e.,

[6] Here $X \ominus Y = (X - Y) \cup (Y - X)$, the symmetric difference between X and Y.

[7] So the set D models your beliefs and determines your consequence operator. This is quite natural if you

always can be approximated using a weaker cumulative (background) semantics. The evidential semantics presented by Seipel ([6]) can be seen as a special case of our framework, taking (partial) minimal model semantics as the background semantics and (partial) stable model semantics as the intended semantics.

References

[1] J.J. Alferes, L. M. Pereira, *Reasoning with Logic Programming*, Springer Verlag, Heidelberg, 1996.

[2] P. Gärdenfors, *Belief Revision*, Cambridge University Press, 1992.

[3] K. Inoue, C. Sakama, Abductive Framework for Nonmonotonic Theory Change. *Proceedings IJCAI'95*, 1995.

[4] D. Makinson, General Patterns in Nonmonotonic Reasoning, in: D.M. Gabbay, C.J. Hogger, J.A. Robinson (eds) *Handbook of Logic in Artificial Intelligence and Logic Programming*, Vol. 3, Nonmonotonic Reasoning, Oxford Science Publications, Oxford, 1994.

[5] S. Marcus, V.S. Subrahmanian, *Relating Database Updates to Nonmonotonic Reasoning* Cornell University, 1994.

[6] D. Seipel. Partial Evidential Stable Models For Disjunctive Deductive Databases. In: T. Przymusinski (ed.) Proc. Workshop on Logic Programming and Knowledge Representation (LPKR'97), 1997.

[7] C. Witteveen and W. van der Hoek, A general framework for revising nonmonotonic theories. In: J. Dix et al., *Logic programming and Nonmonotonic Reasoning*, LNAI 165, Springer, New York, 1997, pp. 258–272.

[8] C. Witteveen and W. van der Hoek, Recovery of Nonmonotonic Theories. Report UU-CS-1997-34, Department of Computer Science, Utrecht University, 1997.